**Networking and Information Technology Research and Development (NITRD) Program: Draft NITRD 2010 Strategic Plan (Request for Public Comment)**

**Dates:** Comments must be received by 5 p.m. EDT on October 11, 2010.

**Summary**: With this notice, the National Coordination Office for Networking and Information Technology Research and Development (NITRD) requests comments from the public regarding the draft 2010 Strategic Plan for the Federal NITRD Program. The draft Strategic Plan is posted at: www.nitrd.gov/DraftStrategicPlan/ Comments of one page or less in length are requested. This request for information will be active from September 10, 2010 to October 11, 2010.

**Address:** Submit comments via e-mail to: nitrd-sp@nitrd.gov
Comments submitted in response to this notice may be made available to the public online or by alternative means. For this reason, please do not include in your comments information of a confidential nature, such as sensitive personal information or proprietary information.

**Overview:** This notice is issued by the National Coordination Office for the Networking and Information Technology Research and Development (NITRD) Program. The draft NITRD Strategic Plan reflects broad input from Federal agencies as well as from researchers and other stakeholders in academia, industry, national laboratories, and professional/technical organizations. Public inputs were solicited in a detailed August 2008 Request for Information (RFI) and in a February 2009 public forum and Webcast. Several hundred comments were received in response to the RFI, and many of these were posted to the NITRD Web site for further comment. The public forum, which included formal presentations by academic and industry experts addressing key concepts for the draft Strategic Plan, was attended by some 100 members of the public, while another 400 persons participated via the Webcast.

**Background:** As required by the High-Performance Computing Act of 1991 (P.L. 102-194), the Next Generation Internet Research Act of 1998 (P.L. 105-305), and the America COMPETES Act of 2007 (P.L. 110-69), NITRD currently provides a framework and mechanisms for coordination among 14 Federal agencies that support advanced IT R&D. These agencies report IT research budgets in the NITRD crosscut, and many other agencies with IT interests also participate informally in NITRD activities. The draft 2010 Strategic Plan for the NITRD Program was developed by the NITRD agencies pursuant to a recommendation of the President's Council of Advisors on Science and Technology (PCAST).

**Invitation to comment:** Inputs of one page or less are welcomed in response to this third and final request for public comment on the Plan. Email to: nitrd-sp@nitrd.gov

**FOR FURTHER INFORMATION, CONTACT:** The National Coordination Office (NCO) at nitrd-sp@nitrd.gov or (703) 292-4873. Individuals who use a telecommunications device for the deaf (TDD) may call the Federal Information Relay Service (FIRS) at 1-800-877-8339 between 8 a.m. and 8 p.m., Eastern time, Monday through Friday.

**U.S. Networking and Information Technology Research and Development (NITRD) Program**

# NITRD 2010 Strategic Plan

# EXECUTIVE SUMMARY

Information technology (IT) – computers, wired and wireless digital networks, electronic data and information, IT devices and systems, and software applications – today provides the indispensable infrastructure for activities across all facets of society. Throughout the IT revolution, the United States has led the world in the invention and applications of these technologies. Ongoing research and development (R&D) to provide advanced IT capabilities for Federal missions has fueled the creation of new ideas, innovators, and innovations addressing key national priorities, including national security, national defense, economic prosperity, scientific discovery, energy and environment, health, individual privacy, and quality of life.

This five-year strategic plan for the Federal government's Networking and Information Technology Research and Development (NITRD) Program responds to recommendations of the President's Council of Advisors on Science and Technology (PCAST). In *Leadership Under Challenge: Information Technology R&D in a Competitive World,* its August 2007 assessment of the NITRD Program, the PCAST recommended that NITRD "develop, maintain, and implement a cohesive strategic plan" that includes "a comprehensive technology vision and strategy that identify the next generation and future generations of important networking and information technology challenges and describe how to meet those challenges."

The NITRD plan presents an overarching vision for the digital world in the 21st century – a world in which "high-speed networks, systems, software, devices, data, and applications are fully secure, safe, reliable, multimodal, and easy to use." In the envisioned future, next-generation IT infrastructure and capabilities will enable continued U.S. leadership in economic innovation and scientific discovery (e.g., alternative energy sources, technologies, and supply systems; personalized bio-genetic medicine; space exploration); national security (e.g., a secure cyberspace; reduced risk of "hot" war); national defense (e.g., dynamic battlefield communications); and education and quality of life (e.g., universal learning technologies and access to information; virtual environments for collaborative work and social interaction; intelligent systems for independent living; life-saving transportation systems; transparent government).

To realize the NITRD vision, the plan calls for advancing U.S. capabilities in three broad areas identified as the essential foundations for sustained leadership in a digital world:

***WeCompute*** – Expanded human-computer partnerships, including more capable, available, and affordable systems; more powerful digital tools for people; and new forms of collaboration between the two.

***Trust and Confidence*** **–** The ability to design and build systems with levels of security, safety, privacy, reliability, predictability, and dependability that "you can bet your life on."

***Cyber Capable*** **–** Transformed education and training to ensure that current generations benefit fully from cyber capabilities and to inspire a diverse, prepared, and highly productive next-generation workforce of cyber innovators.

The plan discusses the topical elements of each foundation and summarizes the principal research and education challenges that need to be addressed, providing a comprehensive research and education strategy for building a bright U.S. future. The plan concludes that the NITRD Program should pursue expanded multiagency collaboration; cultivate new forms of partnership with academia and industry (e.g., in testbeds, large-scale infrastructure such as clouds, data sharing, prototyping, and standards development); and continue to lead by example in multidisciplinary activities and identification of critical-path research needs.

# I. Introduction

## A. Advancing Our National Priorities

Rapid change is often cited as the only constant in networking and information technology (IT). After all, in only a few decades, IT systems have evolved from relatively rudimentary "dinotech" to futuristic electronic marvels that enable virtually all of humanity to communicate, access information, and compete using a device that fits in the palm.

In fact, there has been another constant throughout the IT revolution: U.S. leadership in the invention of digital technologies. As a result of that leadership, IT has significantly enhanced the Nation's economy, quality of life, and national security. Ongoing Federal research and development (R&D) to provide the world's most advanced IT capabilities for U.S. government missions has fueled the creation of new ideas, innovators, and innovations that have enabled the United States to address key national priorities. For example:

- **Economic prosperity:** New multibillion-dollar IT industries have arisen from R&D breakthroughs; communications technologies and standards make possible vibrant e-commerce, more efficient business-to-business interactions, and improved industrial process control systems
- **Quality of life:** Advances in communications standards, electronics, and the Internet have spurred an explosion of novel social networking and web applications with millions of participants, enabling people to collaborate effectively over great distances
- **National security and defense:** Digital capabilities have transformed strategic communications and reduced battlefield risk for military personnel
- **Health and health care:** IT makes possible skilled on-site medical care, in the home and at remote locations
- **Energy and environment:** Sensors and high-fidelity modeling and simulation enable better energy efficiency and weather and climate-change prediction, and speed development of renewable energy sources
- **Education and training:** Learners of all ages now have on-demand access to vast education and knowledge resources
- **Open and transparent government:** Rapidly growing electronic access systems are opening government information and participation to the public and increasing efficiency in government services.

This five-year strategic plan for the Networking and Information Technology Research and Development (NITRD) Program looks beyond the immediate IT landscape toward longer-term transformational impacts of digital technologies (and their yet-to-be-discovered successors) on society and the world. The plan proposes three new concepts that the NITRD agencies view as essential foundations for maintaining U.S. technological preeminence in the dynamic era ahead. The plan also identifies the advances needed in IT research and education that will enable the United States to build these foundations and reap their economic and societal benefits.

# B. NITRD Vision

The NITRD strategic plan begins with a vision of the digital future.

*High-speed networks, systems, software, devices, data, and applications are fully secure, safe, reliable, multimodal, and easy to use. This ultra-high-bandwidth infrastructure – including both flexible, mobile wireless and hard-wired connectivity – is always available, ubiquitous, and can be activated when needed; it is affordably accessible to anyone, anywhere, anytime. Information shared over the infrastructure is completely secure; user privacy and confidentiality can be assured on demand; and pervasive repositories provide for archiving and retrieving data in perpetuity. In the all-encompassing dimension called cyberspace, people – unfettered from constraints of time, circumstance, and location – partner with computing devices and their capabilities to learn, imagine, discover, play, create, invent, interact, and collaborate in real time in ways that enhance life and generate solutions for the world's most complex problems.*

Imagine, for example, that:

*Energy becomes renewable, affordable, and non-polluting, and the U.S. leads globally in production and applications of advanced energy technologies. How?*
* Data-intensive scientific computation combines with human creativity to accelerate discoveries leading to new alternative energy sources, clean production and distribution technologies, and eco-friendly synthetic fuels.
* "Intelligent" distribution networks seamlessly control and manage the integration, flow, storage, and efficient distribution of electricity from all energy sources to "smart" end-user systems that empower consumers to maximize conservation and eliminate energy leaks and waste.
* The new U.S. energy-technologies sectors provide technically skilled jobs in energy production; transportation and power-infrastructure manufacturing and maintenance; and supply chain and services.

*Affordable, high-quality, patient-centric health care is accessible to anyone of any age anywhere and anytime. How?*
* Individuals possess secure, privacy-protected digital health profiles – with physiological and genetic data as well as medical history, treatment contra-indications, and the like – so they can receive personalized medical attention whenever and wherever they need it, while reducing the costs in time and paperwork of providing care; virtual private network capabilities in the home and throughout the community make possible real-time, confidential patient-doctor consultations including examination and diagnostic services that save time and effort for both the patient and the physician.
* The elderly and the chronically ill are able to live independently, supported by smart assistive devices and always available, interactive network connectivity with family, friends, and caregivers.
* Medical errors are reduced because medical devices perform with very high reliability and accuracy, and smart medical networks monitor the systems' functioning and alert personnel instantaneously to anomalies.

* Computational genetic and biomolecular analyses enable researchers to develop techniques, customized to the individual, to prevent the onset of disease or provide cost-effective early detection and personal treatment plans.

*People enjoy increased security, educational and economic opportunity, and quality of life. How?*
* Cyber national security capabilities reduce the risk of cyber warfare affecting major economic sectors and critical infrastructures connected to the cyber infrastructure; digital defense systems reduce battlefield risks and casualties by orders of magnitude from today's levels.
* IT supports a national ecology of innovation that enables us to invent, make, and build novel artifacts that generate new economic opportunities, technology-based industries, and jobs that strengthen U.S. global science and technology leadership.
* High-speed communications and computation infrastructure enables large-scale distributed work environments, reducing the need for commuting and thus urban congestion and pollution.
* The self-managing and "unbreakable" broadband cyber infrastructure enables rapid, efficient processes in every sector, eliminating costly, labor-intensive defensive measures and repairs to networks and systems while enabling advanced emergency-response communications networks whenever and wherever they are needed.
* Intelligent systems help manage accident-free transportation on the ground and in the air, and perform a wide range of physical and intellectual tasks at all scales – from home chores to precision noninvasive surgery to massive ultra-speed calculations and data analyses that enhance human problem-solving and decision-making.
* Teaching, learning, and knowledge environments are virtually universal, making the excitement of discovery an integral part of life from cradle to grave.

## C. The Vision's Three Essential Foundations

The keys to this future of promise will be held by those who successfully develop its three essential foundations. WeCompute, the first foundation highlighted in the NITRD Strategic Plan, identifies the technological direction of the 21st century toward closer, more productive partnerships between people and computing systems, including their data resources. Trust and Confidence identifies fundamental attributes that must be assured in the information technologies upon which society increasingly depends, and Cyber Capable describes the role of education and training in preparing citizens and workers for the cyber future. The NITRD agencies believe all three foundations (summarized below) must be developed to make possible the visionary advances in science and technology that will keep the U.S. at the forefront of economic prosperity, innovation, and scientific leadership.

***WeCompute:*** New understandings and technologies that expand and exploit the intellectual and creative potential of synergy between humans – from individuals to large-scale groups – and computing systems and data. Deepening this dynamic partnership will enable new levels of intelligence, intuition, and awareness emerging as greater than the sum of the parts in the union of cyber, human, and social capabilities. The artifacts we engineer and manufacture will increasingly meld digital, physical, and cognitive attributes in unprecedented ways that enhance the quality of life and extend the frontiers of discovery, innovation, and achievement.

***Trust and Confidence:*** The ability to design and build systems with levels of security, safety, privacy, reliability, predictability, and dependability that "you can bet your life on," and an understanding of "the science of security" that will enable us verifiably to create systems we can trust out of fundamentally untrusted (whether by accident or by design) components. Cyber systems that inspire trust and confidence will greatly increase their value to society, and we will be able to ensure that data resources and systems can be reliably used for their intended purposes.

***Cyber Capable:*** Transformed education and training to ensure that current generations benefit fully from cyber capabilities and to inspire a diverse, prepared, and highly productive next-generation workforce of cyber innovators.

Today, we are a long way from having these foundations for the future in place. Subsequent sections of the strategic plan discuss each foundation in more detail, summarizing the current state of the art and pointing to key areas of research, development, and education in which Federal resources and policy attention will be needed in order to realize the NITRD vision.

# A Digital Future Scenario

In 2020, the United Nations is holding its 4th Earth Summit (ES4), following those in Rio de Janeiro (1992), Kyoto (1997), and Copenhagen (2009). In the previous 10 years, progress in curbing greenhouse gas emissions has been uneven and inadequate, and now many scientists believe, and evidence indicates, that the Earth is approaching a tipping point of rapidly accelerating ecological damage.

In order to minimize environmental impact, ES4 was convened in the Global Community virtual environment, where users attend via headsets that reproduce face-to-face meeting in high-resolution 3-D, provide real-time language translation, and give high-bandwidth access to worldwide IT resources. These headsets allow simultaneous voice, gesture, and tactile input.

The Global Community collaborative space is a secure private enclave in cyberspace that is invisible to global Internet users; set up under a dynamic reservation system, it can be entered only with appropriate user authentication and will disappear once the summit is over. In a hook-up prearranged with conference organizers, the proceedings of the gathering will be watched in real time by hundreds of millions of people around the globe; they will have access to the scientific data emerging, will be able to participate in some distributed computational tasks, provide information, and interact with scientists also watching the event but not directly involved.

Halfway through ES4, carbon reduction negotiations have converged on 3 scenarios. Final agreement is stalled due to lack of agreement of the climatic and economic impacts of these scenarios. However, the delegates agree to adopt impact assessments from an ensemble of world's most respected Earth systems and economic models, if that assessment can be provided within 72 hours. The assessment must have the following characteristics:

- 1, 5, 10-year economic projections
- 10, 50, 100-year climate projections
- Best models and highest resolution feasible
- Interim results daily; consensus results delivered within 72 hours
- On day 2, all nations are to demonstrate reduced carbon emissions, while Smart Grids monitor usage and sensor-nets measure real-time emission.

This will be the largest, most distributed ensemble computation ever attempted, and a milestone in international cooperation. Within a new Global Community virtual environment created for the modeling effort:

- Modeling experts discuss and resolve details about the scenarios
- A multidisciplinary framework is established for efficiently coupling Earth system and economic models, and for determining the consensus of the ensemble of models
- U.S. Federal agencies and countries initiate their surge computing processes

- Citizens organize and communicate about energy-saving initiatives for the 2nd day
- Smart Grids and Earth-observing satellites post measurements during days 1 (normal) and 2 (energy-saving) for inclusion in final modeling and analyses on day 3
- Some computational tasks are farmed out to millions of computers and cell phones of engaged citizens, increasing personal commitment to sustainability
- Data analysis and visualization experts deploy concurrent visualization tools for modeling experts to assess progress, and summit delegates to understand regional impacts, and
- All activities are communicated to the public in real time, yet strong security is maintained to assure that results are valid.

Due in part to NITRD's critical role in coordinating and infusing U.S. IT R&D, the effort is successful, and ES4 concludes with a global consensus on environmental sustainability.

# II. The Challenge of Accelerating IT Change

## *Why We Have Reached a Pivotal Transition*

Computers, wired and wireless digital networks, electronic data and information, IT devices and systems, and software applications form the indispensable infrastructure for activities across all facets of society, driving innovation in every sector. Only a few technologies in human history have had such profound global impacts, and none as rapidly; the digital age has barely begun.

Even so, because the rate of technological change is accelerating exponentially, we are now hurtling into the all-digital future. In 1999, for example, there were fewer than 300 million cell phone subscribers worldwide. Today, there are almost as many cell phone subscriptions as people on the planet, and the volume of wireless Internet traffic is expected to keep soaring as new applications appear. Google statistics tallied a billion unique World Wide Web pages in 2000; in July 2008, the figure passed the one-trillion mark. Today's Federally sponsored Roadrunner and Blue Waters petascale systems can complete in an hour a computation that would have taken more than 100 years for a supercomputer of two decades ago. Other exponential changes – such as the radically shrinking cost per megabyte of storage and per megabit of bandwidth – are speeding the emergence of global social networks, Web services, and the new form of shared distributed computing and storage resources called cloud computing.

Cascading growth in types and quantities of digital devices and applications, along with decreasing cost and size per unit, ultimately will generate fundamental, worldwide changes in the quality of life and the speed of innovation. The NITRD strategic plan illuminates this pivotal time of transition and assesses its implications for U.S. R&D.

# III. Three Essential Foundations of IT R&D

***Where we need to go from here***

The focus of the digital age thus far has been on its novel machinery. Now attention must turn to people and the integrated and seamless use of digital age capabilities and resources. The path-breaking innovations of tomorrow will involve an increasingly intricate fabric of relationships among individuals, organizations, societies, and the technologies providing not only basic information/communication infrastructures and tools for daily life but also advanced capabilities supporting national security, economic innovation, 21st century employment opportunities, accelerated invention and scientific discovery, and individuals' and government's right to shared and private data. The fundamental task before us is to realize the full potential of human-digital partnership in the emerging digital world. The following sections explain how building the three foundations of IT R&D will position the U.S. to transform the challenges of this pivotal transition into strategic opportunity.

## A. WeCompute

Without question, computing systems can perform certain tasks far faster, with greater accuracy and reliability, and at lower cost than people can, and other tasks that are beyond the ability of a person to perform at all, such as:

- Complete in a relative eye-blink a sequence of calculations that a person or a group of people could not complete in their lifetimes
- Process signals, transactions, and records at nearly light speed around the clock and across the seasons without fatigue, including indexing, search, and retrieval for massive data sets in sizes unimaginable in an ink-on-paper world
- Control complex systems such as missile interceptors, telephone switch gear, and automated safety systems that must react to massive, simultaneous inputs at speeds and with levels of precision and reliability greatly exceeding human capabilities

People, on the other hand, can:

- Think creatively – Use imagination and intuition – leaps of faith – as well as abstract reasoning to come up with innovative ideas and inventions
- Grow independently – Continually learn and develop new skills throughout life and apply these in novel combinations for evolving goals and challenges
- Perceive and adapt to nuances – "Read" subtle changes in social circumstances and physical environments and recognize implicit elements to adjust strategy and tactics appropriately
- Be decisive, despite uncertainty – Make decisions in complex situations where there is no single "correct" choice

Deepening the human-digital partnership will require new scientific findings and technical innovations regarding both parties to the partnership – for example, how to engineer computing systems that are not only more capable but also easier for people to work with (e.g., mobile "human-like" systems that interact with people through hearing, sight, speech, and touch; systems that learn and adapt based on their experiences, and assess and maintain their own health). In addition, more precise and detailed knowledge about human cognition and perception

is needed, to enable us to design more intuitive interfaces between people and machines; better computational and data tools to support human brainstorming, collaboration, and efficient decision-making; and more effective applications to improve learning and problem-solving. Enhancing human-computer interactions will also require advances resolving some of the most intractable problems in computer science – e.g., how to engineer software that is at once reliable, robust, secure, cost-effective, sustainable, and easy to use. At the most fundamental level, closer partnership will demand a highly robust global infrastructure with far greater end-to-end speeds, security, flexibility, and capacity than today's Internet.

Such capabilities will make it possible for people to work together and interact with computing in unprecedented ways to achieve results that neither people nor computers alone could attain. Glimpses of the possibilities are already emerging from Federal R&D. For example:

- In 2009, the Air Force for the first time in its history trained more operators of unmanned aerial vehicles (UAVs) than aircraft pilots, and that trend is expected to continue. In April 2010, NASA's unpiloted Global Hawk aircraft made its first atmospheric research flight over the Pacific; the pre-programmed flight, controlled by operators in California, reached altitudes above 60,000 feet – roughly twice as high as a commercial airliner flies; the plane can cover 11,000 nautical miles, or half the Earth's circumference, staying aloft for up to 30 hours.
- In large-scale scientific investigations such as the NITRD agencies' Earth System Modeling Framework activity, agreement on data protocols, standards, and software conventions are enabling many organizations and very large numbers of researchers to share and make use of relevant data sets in real-time collaborations.
- Robotic surgical instruments have moved out of research laboratories and into hospital operating rooms.
- Global volunteer networks of personal computer users are contributing computing cycles to obtain computational results for an array of large-scale research projects, from the study of proteins to identification of objects in the Milky Way.
- The U.S. Department of Transportation (DOT) has unveiled a research plan to revolutionize U.S. motor vehicle transport through an intelligent, multimodal wireless infrastructure and suite of software applications. Called "IntelliDrive$^{SM}$," the aim of the integrated suite of technologies is to improve transportation safety by providing real-time connectivity with and between vehicles, between vehicles and the roadway, and with portable devices that enables all vehicles to sense roadway dangers and communicate them to drivers, along with real-time congestion, accident, weather, and other travel-related information.
- Using high-end computation to analyze the radiation waves under the Sun's surface, a research team from NOAA, the University of Colorado, and the National Solar Observatory has developed a technique to predict with great accuracy the eruption of solar flares. The powerful radiation emitted in solar flares can disrupt Earth's communications systems and endanger astronauts at the International Space Station (ISS). The new capability for the first time provides a 48-to-72-hour advance warning, enabling managers to reposition communications satellites and shut down sensitive systems, and astronauts to vacate the ISS.

Realizing the WeCompute vision will require sustained R&D progress from today's state of the art. Following are the key elements of the vision and the major research challenges to be met in each.

**Making the Digital World Accessible to Everyone (A7)**

In the NITRD vision, the IT infrastructure of the future will be everywhere and always available, making possible **A**nywhere, **A**nytime, **A**ffordable **A**ccess to **A**nything by **A**nyone **A**uthorized (A7). IT capabilities that enable universal participation will radically democratize the IT domain, so that all can contribute to and share in the resources and benefits of cyberspace. At the same time, groups of individuals with similar interests (e.g., a research collaboration) can form and dissolve dynamically to pursue their mutual interests privately and securely as needed.

*Where we are now:* The global Internet points toward the future with its rapid expansion to encompass the burgeoning technologies of wireless networking. But today's Internet is not nearly robust or advanced enough to satisfy the demands of a future in which devices, data, and people at every scale are interconnected and in constant communication, not just worldwide but across outer space. For example, most people on Earth still lack access to the Internet; even in the U.S., one-third of the population currently lacks broadband Internet connectivity at home (lagging 15 other advanced nations), and a majority cites cost and lack of computer skills as key factors. Indeed, access to the information riches and services of cyberspace today remains limited mainly to people who know how to work keyboard-activated devices and have the high-speed network connectivity required to experience bandwidth-intensive Internet applications, such as streaming video and real-time interactivity. (The Administration has announced a National Broadband Plan; its dual aim is to provide jobs by incentivizing companies to both extend broadband connectivity to rural and underserved communities and increase U.S. broadband network speeds, which also lag those of many other countries.)

*Research needs:* In addition to fundamental networking research (see "Evolving and Scaling Socio-Technical Network Infrastructure" below), enabling a more powerful, more scalable IT infrastructure for the future will require advances across the spectrum of information technologies. For example, power consumption must be reduced to enable ubiquitous computation. Language barriers will need to be eliminated through instantaneous language translation. Interoperability issues in data, systems, and software will have to be resolved through agreement on common standards, protocols, and policy regimes; systems must be designed that can adjust to changing environments and the needs of individual users; and substantial improvements in end-to-end performance will be required to provide users with seamless access to resources from their own desktop, laptop, or mobile device. R&D in new materials must be pursued to produce gains in energy efficiency and miniaturization that can continue driving down per-unit costs of IT devices and services, so the IT infrastructure can readily expand to include new participants and uses. R&D to enable specification and enforcement of dynamic security and privacy policies tailored to individuals as well as to organizations will also be a key underpinning of the A7 environment.

**Accelerating the Future of Computing**

Historically, the U.S. has aggressively led advances in high-performance computing (called high-end computing [HEC] in the NITRD Program) because HEC provides a competitive national advantage, supports science and technology leadership, and plays a critical role in advancing Federal missions and other national priorities. The HEC technical area encompasses all of the challenges of leading this advancement, including not just better (e.g., massively scaled-up, more efficient, and resilient) system hardware and software, but also more efficient provisioning

models (e.g., surge and cloud computing), improved mathematical and computer science underpinnings for analysis and modeling of multi-scale and ultra-scale data, and new programming environments and tools for easier development and usability of advanced scientific applications.

*Where we are now:* The Top500 supercomputers list, which has chronicled a decades-long exponential advance in HEC technology, shows China and others now nipping at our heels. The Top500 analysis also extrapolates future HEC performance, indicating that we will continue to see 1,000x increases in capability every decade or so. The most powerful Federal leadership systems have achieved petascale speeds (1,000 trillion floating-point calculations [flops] per second) and are expected to reach the exascale (a quintillion flops, or $10^{18}$) within a decade. At the same time, however, our current means of reaching these ever-higher computational speeds – packing multiple processors into every computer chip – is creating a crisis in our ability to program system and application software to efficiently exploit the emerging many-core and heterogeneous computing architectures. Moreover, as the number of cores and components rises, the likelihood diminishes of correctly completing data-intensive computations of increasing scale and complexity. In addition, the cost to power large HPC facilities is growing unsustainably rapidly.

*Research needs*: The massive parallelism necessary to enable software to fully exploit the speeds and computational capabilities of supercomputing systems with heterogeneous components and up to millions of microprocessors presents enormous challenges for both system and application designers. Developing robust code that can be partitioned into many parallel subroutines for efficient multicore/many-core processing and then reintegrating the results as final output require breakthrough advances in the underlying mathematics, design, and engineering of HEC software. The goal of making HEC environments easier to use and more productive, reducing time to solution, remains elusive; even at today's levels of system complexity, down time due to software issues is rising as a proportion of total operational costs. One promising R&D avenue is resilience – concepts and technologies enabling a HEC system to continue to function amid software faults, anomalies, and errors, and to alert operators about problems without necessitating a shutdown.

Also critical to the long-term future of U.S. supercomputing leadership is research in technological approaches to reduce the steadily rising energy demands of large-scale HEC systems and facilities, which typically consume many megawatts per year for operations and cooling. Because advances to change this unsustainable energy-use trajectory may arise from multiple research fields, the search for scientific breakthroughs must be pursued across the board, in power management and heat dissipation technologies, new materials such as nanoscale composites, novel power-saving platform and system architectures, computing technologies (e.g., nonvolatile computing), and computational methods (e.g., spintronics, analog computing), as well as in next-generation computing concepts such as quantum information science. Advances in nano, biological, and quantum sciences also may lead the way to radically different system architectures and computational methods, providing the basis for next-generation leadership in computing at all scales.

**Evolving and Scaling Socio-Technical Network Infrastructure**

Network connectivity, beginning at the computing platform or device and radiating outward through local-area networks connecting to wide-area networks connecting to the Internet's network of networks, creates the wired and wireless communications fabric that makes a digital world possible. Over the last four decades, Federal investments in basic network research have led the way to the Internet, the World Wide Web, wireless mobile and optical networking, and an array of network-based applications that are reshaping societies and economies around the globe.

*Where we are now:* Advanced computer networks provide the infrastructure for transporting, developing, archiving, accessing, and using the huge volumes of data that support critical functions in every sector, such as storage and nearly instantaneous interchange of data in the financial markets. Scientific experiments such as the Large Hadron Collider at CERN distribute petabytes of data to thousands of scientists around the world who are seeking to uncover the fundamental nature of matter and the "dark energy" that dominates the universe. Dynamic, heterogeneous, secure, and reliable networks are also critical to DoD's ability to defeat adversaries, to DHS's ability to respond to natural disasters and terrorism, and to Federal efforts to improve health care for all.

Like the LHC, a growing number of scientific applications require very large bandwidths to support massive data transfers and the need for near-real-time coordination and data transmission protocols tailored to the data requirements. Other such applications include Earth system modeling supported by the Earth Systems Grid (ESG); computational genomics; and Very Long Baseline Interferometry (VLBI), a radio astronomy application enabling simultaneous observations of an object by many telescopes combined, emulating a telescope with a size equal to the maximum separation between the telescopes.

Current high-performance networking (including science networks) does not generally support these demanding requirements; the current approach is to provide dedicated point-to-point network links and services among the key researchers (or to have researchers move to sites where adequate networking is available). Today, such services lie mostly outside the science networking provided to the larger research community and are implemented outside the university science networking infrastructure.

*Research needs:* The network infrastructure of the 21st century must be made robust enough to meet very diverse demands, including network services supporting A7; exponential increases in data volumes and changes in how people access data (e.g., data in the network); ultra-reliable, secure networking (e.g., for national security and the commercial and banking sectors); new networking technologies that scale (e.g., all-optical networking) to provide the end-to-end bandwidth, performance, and services required for data-intensive science; and heterogeneous networking (e.g., wireless, optical networking, satellite communications, and others). The following are core areas of networking R&D in which advances are needed:

- **Foundations:** Architectural principles, frameworks, and network models to deal with complexity; heterogeneity; multi-domain federation, management, and transparency; end-to-end performance; and differentiated services.
- **Design:** Secure, near-real-time, flexible, adaptive services with built-in intelligence to facilitate discovery, federation, and management of resources across domains and to increase

the application robustness and resistance to attack even in extraordinarily complex systems and new ways of interconnecting networks to provide those services.

- **Management:** Management methods and tools that incorporate intelligence in the network to enable effective deployment, control, and utilization of complex networks and resources in dynamic environments, across domains, and with ever-increasing network and application complexity.
- **Privacy and Security:** Achievement of a high degree of security even in complex, heterogeneous federation and policy environments, especially in the face of component failures, untrusted components, malicious activities, and attacks, while also respecting privacy and maintaining usability, e.g., provide scalable federated policies for authentication, authorization and accountancy
- **Usability:** Adaptable, user-centered services and interfaces that promote efficiency, effectiveness, and fulfillment of user needs without overwhelming users with unnecessary or unauthorized data.

This agenda must be pursued across the spectrum from fundamental to applied research and with engagement of all sectors to attain widely deployable innovations. Necessary elements include:

- Basic and applied research in the full range of network architectures, theoretical models, analysis techniques, hardware, software, security and privacy, and middleware needed to support the next generation of uses for networks and explore new paths to develop capabilities that cannot be supported on the current evolutionary path
- Partnerships with application developers, users, and stakeholders to test basic research ideas on real problems in areas including national security, support of scientific leadership, and human health
- A suite of testbeds and prototype networks that enable understanding and creation of new technologies, data systems, and improvements in end-to-end performance at varying scales. The massive size of existing deployed networks such as the Internet limits research and development, while laboratory and simulation studies cannot address some aspects of the solutions, particularly complexity, their ability to scale, and their potential realism. The testbeds and prototypes will range from high-flexibility/low-cost platforms to high-performance embedded systems.

Research, partnerships, and testbeds and prototype networks are closely interrelated. Testbeds and prototypes are needed to test and develop new networking capabilities in realistic environments, to assure they can be implemented technically and economically, and to explore policy frameworks. Partnerships between the researchers and the application developers will help assure that R&D capabilities can be implemented in real networks and that other application resources such as computing and storage are provided.

**Creating the Smart Planet**

A smarter world will be one in which all kinds of objects, devices, and large-scale physical systems are interconnected, compute-empowered, and instrumented to perform tasks (monitoring, regulating, measuring, analyzing, alerting, etc.) with, on behalf of, and in the best interest of people. This infrastructure will also enable people to collaborate in real time, dynamically creating short-term ad hoc networks linking them to devices, data and information, computing platforms, and applications as needed. Some components of the smart planet

infrastructure will be stand-alone robotic systems designed to perform tasks autonomously; others will be what are now called cyber-physical systems. These are networked computing systems – interconnected software, microprocessors, sensors, and actuators – deeply integrated within engineered physical systems to monitor and control capabilities and behaviors of the physical system as a whole. Such systems are already essential to the effective operation of U.S. defense and intelligence systems and critical infrastructures (e.g., air-traffic-control, power-grid, and water-supply systems), industrial-process control systems, and other large-scale civilian systems, as well as to smaller-scale applications in cars and medical devices. Demand for and uses of cyber-physical systems are growing worldwide.

*Where we are now:* Federal investment in embedded computing, networking, and control has been relatively limited over the past two decades. Computing was once a minimal component of engineered systems, and systems were designed to be operated separately and in benign or controlled environments. Now the "cyber" aspects of engineered systems and products are becoming the very key to making these systems more capable. And the need for deployment is increasingly in situations where systems must be designed to interact and cooperate, often with high degrees of autonomy. This is illustrated in the rapidly growing demand for increased capability in transportation (e.g., safe routing, collision avoidance), manufacturing (precision control, using new – even cyber-physical – materials), agriculture and mining (robotics), and medical diagnosis and therapies (implanted sensors and actuators). Society benefits when surgery can become less invasive, reducing recovery times. Advances in computer-controlled robotic and laser surgery (such as those enabled by the daVinci medical robotic system) are in this direction. The U.S. industrial economy has depended upon the productivity of its workforce – enabled by its technological capability – for relative U.S. strength in industrial sectors worldwide. That lead has declined as other nations have rapidly joined the technology race and have sought to produce ever more sophisticated products and systems.

*Research needs:* The next Industrial Revolution will be one of cyber-physical systems. A new systems science is needed to provide unified foundations, models and tools, system capabilities, and architectures that enable innovation in highly dependable cyber-enabled engineered and natural systems. Better understanding of system complexity is also necessary in this research area to aid in improved management and decision support. Specific technical areas for emphasis include:

- Unifying foundations for modeling, predicting, and controlling systems that exhibit combined cyber (logical/discrete/digital) and physical (continuous/analog) system behaviors
- New approaches for supervisory control of systems that must interact on an ad hoc basis.
- Scientific and engineering principles, metrics, and standards that integrate the disciplines of real-time embedded systems, control, communications/networking, security, human-machine interaction
- Technology to close the design and productivity gap between modeling, programming, and runtime execution of cyber-physical systems
- Principles for reasoning about and actively managing properties of complex, multiscale, real-time cyber-physical system interactions, including: safety, security, reliability, performance.
- Design methods and systems technology for autonomy, human interaction, and management of control authority
- Open systems approaches for composition, integration, and coordination of cyber-physical systems

**Enabling Complex and Sophisticated Software**

Like networking technologies, software makes the digital world possible, directing the functioning of computers and devices and providing the electronic instructions for the applications of computing that shape our lives. For example, NASA spaceflight software controls a preponderance of overall system functionality. A complex software system can be defined as a system comprising interacting "simple" software modules that, working together, exhibit a high degree of complexity resulting in a higher-order behavior. The more complex the software system is, however, the greater the chance for unpredictable emergent behavior, which increases the risk of system failure with potentially significant impacts to the businesses, services, equipment, or users depending on the systems.

*Where we are now:* Critical U.S. defense, security, health care, and economic capabilities depend on complex software-based systems that must remain operational, useful, and relevant for decades. Today's software design and development tools and practices can make any of these goals difficult to achieve. For example, consider keeping software relevant for decades: the requirements originally used to design the software often change multiple times during the development phase, then many more times during the continued use of the software system. How can the need to keep the software relevant and useful be balanced with the need for software that is well defined, tested, and meets evolving operational requirements? The persistent and widening gap between the quality of hardware and that of software continues to burden systems development and broader efforts to innovate in networking and information technologies.

*Research needs:* The tradition of incremental changes in software development provides an inadequate basis to address the complexity of contemporary critical systems. Improving the quality and cost-effectiveness of this software constitutes a core technical challenge that requires breakthrough innovations, ranging from the fundamental science and engineering of software to the application level. Research is needed to rethink software design – from the basic concepts of design, evolution, and adaptation to advanced systems that seamlessly integrate human and computational capabilities. New practices, technologies, tools, and measurement methods are required that can reduce the errors, defects, and vulnerabilities that occur during software development. Specific research topics include:

- Foundational principles for software design
- Formalized science-based software architectures and design methods
- Tools and principles to build, maintain, and expand ultra-large-scale software systems
- Programming languages, tools, and practices for modeling, designing, developing, testing, and validating software
- Tools and practices for improving the interoperability and usability of software applications
- Repositories of software design and development knowledge and reference software
- Improved software assurance that reduces or eliminates software defects, weaknesses, and vulnerabilities through improvements in automated test methods, measurement methods, technology and tools, and guidance and standards for development of trustworthy systems
- Parallel programming languages, compilers, operating systems, environments, and models
- Software for computation- and data-intensive applications
- Software effectiveness metrics
- Highly user-friendly and interactive software systems

**From Data to New Knowledge**

Most of the world's information is now "born digital," and legacy texts, images, sounds, videos, and films as well are being digitized around the clock. Although statistical estimates vary, they agree that the amount of digital data generated annually is many orders of magnitude greater than the total amount of information in all the books ever written, and the total is expected to continue growing exponentially. In the advanced sciences alone, the proliferation of ultra-powerful and distributed data-collection instruments and experimental facilities has turned the conduct of leading-edge research into a global-scale, data-intensive enterprise. The Federal agencies in the NITRD Program together generate exabytes of research data annually. Financial, commercial, communications, and Web-based enterprises likewise generate vast amounts of new digital information on a moment-by-moment basis.

*Where we are now:* Today, our capacity to create electronic data is outpacing advances in the technologies needed to manage and make effective use of society's data resources. Ultra-large-scale data sets – what scientists refer to as "big data" – are troves of potential new knowledge, but as noted above, the current networking infrastructure does not provide levels of end-to-end performance that would enable individuals and groups to access and work with big data on their desktops. While the plummeting cost of mass storage eases the stress of archiving massive data resources, we also do not yet know how to design scalable technologies for rapidly identifying, integrating, refining, analyzing, and visualizing heterogeneous and ultra-scale information in ways that help people learn, think, and decide. Nor do we yet have a rationalized, robust information infrastructure for the long-term preservation, curation, federation, sustainability, accessibility, and survivability of vital Federal electronic records and data collections, such as those overseen by NARA. *Harnessing the Power of Digital Data for Science and Society,* the 2009 report of the Interagency Working Group on Digital Data (which includes many NITRD agencies), has proposed an initial framework for developing such an infrastructure.

*Research needs:* We need far more powerful and nuanced tools than exist today to mine data troves deeply, and to combine diverse forms of data, in order to find significant items, patterns, and relationships that could lead to new insights. To support complex human, societal, and organizational ideas, analysis, and timely action and decision-making, multisource forms of large-scale, raw digital information (e.g., sensor data) must be managed, assimilated, and accessible in formats responsive to the user's needs and expertise. At the extreme scale represented by 21st century scientific and other data, significant R&D challenges in applying information to enhance discovery and decision-making remain to be addressed, including:

- **Information standards:** Data interoperability and integration of distributed data; generalizable ontologies; data format description language (DFDL) for electronic records and data; data structure research for complex digital objects; interoperability standards for semantically understood ubiquitous health information records; and information services for cloud-based systems
- **Decision support:** Next-generation machine learning and data mining algorithms; portals and frameworks for data and processes; tools for large-scale collaboration; user-oriented and collaborative techniques and tools for thematic discovery, synthesis, data provenance, analysis, and visualization for decision making; mobile, distributed information for emergency personnel; management of human responses to data; collaborative information

triage; portfolio analysis; development of data corpora for impact assessment and other metrics of scientific R&D; and multidisciplinary R&D in ways to convert data into knowledge and discovery

- **Information management:** Intelligent rule-based data management; increasing access to and cost-effective integration and maintenance of complex collections of heterogeneous data; innovative architectures for data-intensive and power-aware computing; scalable technologies; integration of policies (differential sensitivity, security, user authentication) with data; integrated data repositories and computing grids; testbeds; sustainability and validation of complex models; and grid-enabled visualization for petascale collections

**Pioneering Socially Intelligent Systems**

Emerging types of collaboration, communication, and cooperation – from open source software development, crowdsourcing, and clickworking – illustrate the potential for a new form of "social intelligence" that melds human and computer abilities. Socially intelligent systems can range in scale from one person and one machine to many people and many machines distributed over the globe. WeCompute envisions enhanced environments and tools for large-scale collaborative problem-solving in which the intelligence of large numbers of people, operating in real-time computational environments, can accelerate solutions to the most complex problems by simultaneously letting humans do what they do best (e.g., deriving meanings from sensory inputs, synthesizing disparate experiences, drawing inferences) and machines do what they do best (e.g., fast, complex computation). Such environments would include the far more powerful analytical tools described above to enable people to make effective decisions.

The "intelligence" exhibited in these systems will mimic human capability to reason, perceive the environment, and collaborate with humans and other machines. An intelligent system will learn from all past experience and adapt over time as well as understand people's cognitive and social abilities. The "social" aspects of the system will result from optimizing human actions, interactions, knowledge, and skills in relation to overall goals. Socially intelligent systems may be designed to act autonomously so that humans can remain "out of the loop," as in dynamic allocation of bandwidth or recovery of transportation systems during emergencies. But human-computer partnerships allowing for new forms of complementary engagement may be the most effective of all.

*Where we are now:* Online commerce and communication have been revolutionized through innovations such as micro-blogging, security, video, recommender and reputation systems, and science has been advanced through cyber-enabled discovery and virtual organizations. Even so, today's systems are merely suggestive of the powerful versions of social intelligence that may be developed in the future. The goal is to advance knowledge at the frontiers of computationally mediated human-machine interactions that would reframe the meaning of intelligence.

*Research needs:* New findings about how the mind perceives, evaluates, categorizes, synthesizes, analyzes, retains, and makes decisions about inputs – both internal and from the outside world – will provide researchers with models for engineering intelligence into the capabilities of networking and computing technologies. Research needs to focus on computers as intelligent participants whose perception, computational capabilities, and learning may be unique and able to scale at rates difficult for humans to grasp. At the same time, we also need a better understanding of how people best coordinate, collaborate, and participate in collective action at

such large scales and in real time. One goal is to better understand the types of problems that are best suited for these types of human-computer partnerships. Key areas of research include: machine learning and artificial intelligence; immersive environments and 4-D touch displays; understanding, modeling, and managing complex systems; computational photography; graphics and visualization; social and humanoid robotics; speech recognition, natural language understanding and dialogue systems; and mechanization of economic theories (e.g., n-way kidney exchanges).

Some research questions include:
What is the nature of this collective intelligence? Can we build computational models of political discourse to better understand each other's point of view and resolve disputes? Can a distributed collection of machines and people learn to collect data, analyze them, and ask good research questions, ultimately changing the way science is conducted? How do we best understand the human capabilities that outperform computers and then harness those assets in new human-computer partnerships? How are value systems (i.e., cultural, ethical, legal, etc.) embedded in algorithms and collective enterprises and how should they be evaluated? What new design techniques and methods would result in a broad array of behaviors and achievements that could effectively address current social issues (e.g., emergency response)? How can human needs and values be strengthened through socially intelligent systems? What new theories could explain the behaviors of these complex, dynamic systems?

# B. Trust and Confidence

A screen freeze may be a trivial inconvenience, but the consequences of more serious IT flaws – such as a digitally controlled diagnostic system whose malfunctioning delivers lethal doses of radiation, trains that have a fatal head-on collision due to a software error, or the theft of sensitive information over the Internet – undermine society's trust in the efficacy of IT as the basic infrastructure of the digital age. The perspective of the NITRD agencies is that one of the most significant tests of technological leadership in the years ahead will be the ability to engineer and build IT systems that inspire high levels of confidence because they function as intended safely, securely, reliably, and cost-effectively. Fundamental research to ensure that digital networks, systems, devices, applications, and communications processes earn and deserve the trust and confidence of society thus constitutes an essential foundation for the Nation's future. Since technology is only half of the equation, this work should include a robust interdisciplinary R&D agenda in the behavioral, ethical, legal, and societal aspects of achieving trust and confidence – for example, how to make systems much more user-friendly, so it is easy for users to "do the right thing" in engaging security and privacy-protection features. Following are the key elements of this vision and the major research challenges to be met in each.

**Making the Digital World More Trustworthy**

The necessity for trust and confidence spans far more than the interconnected networks, systems, and software of the Internet and the information residing in those systems. It encompasses the networked computing systems that are deeply integrated within complex life- and safety-critical physical structures such as power grids, buildings, airplanes and spacecraft, ground transportation, and medical devices; and it includes stand-alone computing systems that also perform critical tasks on which human life, safety, and security depend.

*Where we are now:* Over the past decade, we have become increasingly aware as a society of the vulnerabilities associated with our IT systems and infrastructure. The reality is that many of these technologies were invented and engineered before the security implications of pervasive societal reliance on IT systems and networks came to the fore. In the national security, aviation, and space exploration arenas, Federal research has long pursued technical means of assuring that networks and systems can continue to function in adverse environments and amid internal faults and failures; but to date, system redundancy remains the principal failsafe. Since 9/11, Federal agencies, in partnership with private-sector stakeholders, have also focused on research to harden against cyber attack the process-control systems of critical U.S. infrastructures that rely on Internet connectivity. In broad terms, however, efforts to increase IT reliability, safety, and security continue to target add-on fixes for existing technologies rather than new concepts, designs, architectures, and security standards that would incorporate those attributes from the ground up.

*Research needs:* Evolutionary system hardening and software patching will continue to be necessary in dealing with the legacy systems of prior decades still in service. Only foundational basic research, however, can produce the advances needed to make possible inherently more stable, reliable, safe, secure, self-diagnosing, self-healing – and thus far more cost-effective – systems, software, and devices for the dynamic environments of a fully digital world. A fundamental science of security must be developed as an essential component of high-quality IT design and engineering across all technologies and application domains. The science of security must also infuse curricula and training activities at every educational level. Multiple dimensions of the security challenge are described below.

**Securing Life in Cyberspace**

As the President's May 2009 *Cyberspace Policy Review* notes, the Internet's global fabric of near-instantaneous interconnectivity is at once transformative and fragile – beset by the unintended consequences of its multi-decade growth and survival in increasingly dangerous times.

*Where we are now:* The vast sea of information that flows over the Internet and is stored in Internet-connected systems mostly is not secure, nor are the networks and systems themselves. The basic openness and anonymity built into the Internet's trust-based legacy architecture – combined with a seemingly endless assortment of hardware and software vulnerabilities in computing systems – are exploited around the clock by hackers, criminals, and U.S. adversaries. According to some experts, the networks of zombie attack computers called "botnets" today constitute the largest supercomputer in the world. The lack of end-to-end security in cyberspace costs organizations in all sectors many billions of dollars annually; it also threatens major U.S. government objectives, such as reforming the health care system with the aid of health IT and stimulating economic innovation. Further, the interconnections of the Internet with critical infrastructures and systems (e.g., financial) provide vectors for potentially devastating cyber attacks. Currently, attackers have the upper hand (anonymity; stealth; rapidly shifting and increasingly damaging methods; asymmetric strength); defenders rely for the most part on a never-ending cycle of patching networks and systems, but this defends only against previously identified threats, not the constantly emerging new ones.

The Federal government has initiated high-priority efforts to improve coordination of cybersecurity R&D across Federal agencies, with the goals of better securing government information and networks and expanding collaboration with the private sector to address cybersecurity objectives. Because much of the digital infrastructure lies in the private sector, however, developing R&D partnerships and technology deployment strategies acceptable across sectors outside the Government presents complex challenges.

*Research Needs:* The goal of cybersecurity R&D must be to provide *end-to-end* security in networked environments. The immense dynamism and complexity of global networking make this goal a grand challenge for which there will be no single solution. Advances of many kinds are needed, in the policy and educational arenas as well as in diverse technologies. In addition to more inherently secure components, new methods for proactive approaches to improving cybersecurity must be pursued, such as dynamic security; stronger global-scale identity management; better situational awareness; new means of attack attribution and combating malware, botnets, and insider threats; enterprise-level security metrics for assessing the relative effectiveness of policies and techniques; cybersecurity education; and easy-to-use security techniques. One conceptual approach being advanced by the Federal cybersecurity community specifically focuses on ways to eliminate the cyber attacker's advantage over the defender – for example, by employing dynamic virtualization to make attack targets much harder to pinpoint or by creating "tailored trustworthy spaces" on the Internet that provide elevated levels of security and privacy.

**Systems You Can Bet Your Life On**

Cyber-enabled engineered systems in which cyber capability is deeply embedded at all scales must remain safe, secure, and dependable – i.e., "systems you can bet your life on." The challenge of building such systems, often from fundamentally untrusted components, spans essentially every engineering domain. It requires the integration of knowledge and engineering principles across many computational and engineering research disciplines (computing, networking, control, human interaction, and learning theory, as well as electrical, mechanical, chemical, biomedical, nano-bioengineering, and other engineering disciplines) to develop a "new CPS system science."

*Where we are now:* The complexity of cyber-physical systems, including robotic and autonomous systems, is at a point where current methods are inadequate to anticipate possible failure modes and guarantee safe, predictable, efficient operation. The world's leading automotive manufacturers, for example, have suffered recent catastrophic product failures that resulted in enormous recall costs and loss of consumer confidence with potentially huge economic consequences. Deaths due to infusion pump failures have reached such a high level that the FDA has found it necessary to mount an initiative to investigate. At the same time, our appetite for systems of this kind – in which the engineered systems are monitored and controlled by computer and communication networks – drives the steep upward growth curve in system complexity. The dynamic, often decentralized nature of these complex systems places unprecedented demands on the contributing areas of real-time computing, communication, networked control; on engineering for the physical domains; and on verification, validation, and certification support for all of these. Over-design currently is the only path to safety and successful system certification, leading to a mindset of optimizing for a narrow task instead of encouraging adaptability and evolvability.

*Research needs:* We need to establish new, unified scientific and engineering foundations to securely, safely, and systematically understand, build, manage, and adapt these "complex" systems – cyber-physical systems that remain reliable as they interact across internal subsystems, with each other, with human users, and with highly complex and uncertain physical environments – and we need successful exemplars of such systems. A core element of this agenda is development of new, more cost-effective approaches to certifying the quality of these systems, a challenge that today consumes, for example, an estimated 50% of the resources required to develop new, safety-critical systems in the aviation industry. Essential R&D areas include:

- Comprehensive integrated design approaches for cyber and physical system events and actions – for example, exploration and simplification of both nominal and failure mode design for complex system environments that may incur undesirable or hazardous emergent behavior.
- Improved models of system and human behavior that can provide a framework for human-system and system-system interoperation that can enforce safe operation in mixed-initiative systems. For example, models should support interaction without introducing problems such as mode confusion or technology surprise.
- New approaches to fault tolerance: system prediction, recovery, and adaptation technology for rapidly identifying and avoiding potentially reachable failure states, and for maximizing effective fall-back and fail-safe recovery when these cannot be avoided.
- Hardware, software, and control platforms and frameworks that support rigorous – checked or verified – composition of system components and guaranteed regulation of component interactions.
- New scientific underpinnings and design approaches for securing cyber-physical systems, addressing both cyber and natural or malicious physical disruptions (and interactions of these).
- Technology-supported certification approaches that are based on claims and analytic evidence, rather than merely process-based checklists, and that can support modular certification of components and assemblies, with incremental re-certification after system modification – e.g., knowing and certifying what exactly it is that you can trust, and why you can trust it.
- A new generation of design and analysis toolchains that can produce rigorous evidence for high-confidence system design, implementation and evaluation. For example, these would integrate discrete and continuous mathematical models and support rigorous reasoning about the interacting behaviors of cyber and physical components. They might include frameworks that equally support evaluation, V&V, and certification activities, in addition to design and implementation.

**Information Assurance and Sharing**

A primary function of cyber infrastructure is to provide for the safe, secure creation, transmission, storage, and retrieval of all kinds of digital information – including sensitive data belonging to individuals, private-sector organizations, and government. Ideally, both the creator and any recipients or viewers of nonpublic digital information should be authorized and should be able to access it securely; identify its origin and history, or provenance; authenticate its integrity (no one has tampered with the content); and maintain its confidentiality as required. The

information assurance field looks at cybersecurity specifically from the perspective of what is required to maintain the confidentiality, integrity, and availability (CIA) of data.

*Where we are now:* Current concerns in information assurance range from protecting the bits themselves through guarding the larger digital environments in which they reside. Technical areas include security governance and privacy policies; network and system access controls (administrative, logical, physical), identity authentication, management, and non-repudiation technologies and policies; cryptographic techniques for data encryption and decryption; and forensic capabilities for identifying security breaches. In highly sensitive information environments such as DoD, "mission assurance" also employs risk analysis and management tools to analyze and mitigate the security risks of environments in which information is shared across multiple security levels. Today, encryption techniques provide the only data-based direct means of preventing unauthorized persons from obtaining access to digital data. Public Key Infrastructure (PKI) exemplifies the approach of creating a trusted multi-domain network environment, but PKI has been slow to achieve widespread adoption because it is costly and cumbersome to administer.

*Research needs:* The same characteristics of complex enterprises that enable network and information managers to institute access controls and security monitoring – large-scale system homogeneity, static configuration, and software monoculture – also make it easier for cyber attackers to access, tamper with, or destroy information. To realize the NITRD vision, research must seek fundamental advances in hardware, software, and network architectures that can provide immunity from tampering and attacks, possibly by identifying and actively defeating them or increasing system diversity. Likewise, next-generation approaches are needed for securing digital information itself. Exploration of such techniques as homomorphic encryption, for example, may lead the way to data formats that are intrinsically encoded but still usable in controlled environments.

**Understanding the Trade-offs: Balancing Security and Privacy With Other Values**

Designing a system or a network that satisfies a single design goal – security – is still a grand challenge in computer science research. Yet, in reality, many systems and networks like the Internet that are used by real people have to satisfy not just one goal but an array of them. For example, they need to be usable; they should give users the information and personal privacy they expect; they should be open enough so that users can connect with others at a distance and obtain information that is available on other systems and networks. At the same time, the systems need to provide the level of security required by the end users. In many cases, however, it is not possible to satisfy all competing goals.

*Where we are now:* We are just beginning to come to grips with the implications of such conflicts. If it is not possible to design a system that is simultaneously secure, privacy-preserving, usable, and open, what are the potential trade-offs among those attributes? We need to better understand these trade-offs and expand the space of possible solutions.

*Research needs:* Research is necessary to investigate how we can more effectively comprehend, model, and optimize an array of conflicting design goals. In addition, new tools are needed to enable IT managers and end users both to monitor and act to mitigate security risks.

# C. Cyber Capable

The third essential foundation for the bright future envisioned by the NITRD agencies involves the human side of the human-computer partnership. To remain at the forefront, the Nation will need a cyber-capable citizenry – Americans well prepared to be savvy consumers and users of IT and, with advanced education and training, to generate the discoveries, inventions, and creative ideas that will propel U.S. innovation in an increasingly competitive global economy. In the NITRD vision, cyber capabilities and tools will be applied to provide learners at every level with rich and exciting knowledge environments informed by the science of learning, tailored to individual needs, and guided by educators with theoretical and practical IT knowledge. The cyber-capable society will thus require an education and training system that can deliver world-class preparation in computer science and other science disciplines, technology, engineering, and mathematics (C-STEM).

Following are the key elements of this vision and the major research challenges to be met in each.

**A Workforce of Cyber Innovators**

Tomorrow's workforce will have to be agile, adaptable, well educated and trained, and able to keep learning continuously to take maximum advantage of technological advances and contribute to American innovation. This applies not only to cyber professionals, who even today struggle to stay current with their rapidly changing and advancing field, but to professional and technical workers in every sector.

*Where we are now:* A 2009 study conducted for the NITRD Program notes that two IT-related occupations – network systems and data communications analyst, and computer applications software engineer – are among the five fastest-growing in the U.S. economy, and the only two of the five to require a college degree. According to BLS projections reported in the study, the professional and technical workforce in networking and computing should expand by more than 1.2 million, or 24 percent, to 3.5 million between 2006 and 2016. The professional/technical workforce over all is expected to grow by 17 percent over the same period. Projections that include IT-related jobs that do not necessarily require a college degree (such help desk specialists, electronic records processors, tellers, etc) double the size of the IT workforce in 2016. By contrast, the number of computer science and electrical engineering degrees at all levels has been declining since 2004, as has the percentage of degree holders who are U.S. citizens or residents. Government and private-sector employers alike report difficulty finding people with the requisite IT skills.

Labor market projections for the IT workforce, however, do not capture the reality that a very broad range of occupations increasingly involves applications that require IT knowledge and skills. Nor can statistical projections serve as a guide for assessing the adequacy of the educational system to prepare a workforce that leads the world in advanced innovation.

*Research and education needs:* Information technologies are interdependent and are developed from an inherently multidisciplinary basis in the sciences and in engineering. Building systems

and large-scale applications takes teamwork across diverse technologies and academic fields. Moreover, IT capabilities are used in a wide variety of social contexts that IT professionals also need to understand in order to create and use applications effectively. For example, in the 1990's the lack of professionals trained in both computer science and biology prompted NIH to establish the Nation's first graduate fellowship programs in bio-informatics; as a result, such training is now part of the curriculum at many graduate and medical schools.

The PCAST argued in its 2007 NITRD review that the traditional disciplinary stovepipes of the formal educational system present a substantial barrier to development of diversified, broadly interdisciplinary new generations of cyber innovators. We need advances in thinking about how to organize education and training curricula and experiences, particularly at the postsecondary level, to help students develop the intellectual capacity to synthesize knowledge from multiple disciplines and work collaboratively on complex interdisciplinary problems, whether the setting is IT for advanced manufacturing or for a regional social services delivery system.

**The Education of Cyber-Capable Citizens**

Our society will benefit when all K-12 students gain a better understanding of how digital technologies work and how to use their applications safely and wisely. Innovation in several dimensions of education could accelerate this process: employing learning technologies in all grades and subjects; incorporating "computational thinking" – the concepts, mathematics, and logic of digital processes – throughout the formal curriculum at all levels; and expanding outreach efforts to raise public awareness and better inform people of all ages about best practices for IT users.

*Where we are now:* Early IT educational applications offered mainly reading lessons or "skill and drill" testing; emerging science-based knowledge about learning (e.g., its neural basis, psychological theories of knowing, and biologically inspired learning algorithms) is informing the development of more sophisticated learning technologies. In addition, precollegiate educators could exploit the types of computational resources that have transformed the conduct of science and engineering – e.g., authentic and realistic data, digital telescopes, immersive environments, mobile and portable devices, modeling and simulation capabilities, sensor networks, remote instruments – to transform classroom learning. Deploying such approaches widely will require educators able to apply IT expertise to subject-matter pedagogy.

Today, however, according to computer science (CS) experts who spoke at a NITRD strategic planning public forum in 2008, the K-12 curriculum in computer science is extremely limited, mainly focused on beginning programming at the senior-high level. The Computer Science Teachers Association has reported that the proportion of high schools offering an introductory CS course dropped from 78 percent in 2005 to 65 percent in 2009; only 27 percent offered an Advanced Placement (AP) course, and about 11 percent of AP test takers were CS students. As *The Washington Post* noted, "It would be hard to find a student who has never used the Internet for a research assignment or played a video game, but few know much about how computers and the Web actually work."

Recognizing the implications of the disconnect between these trends and the requirements of the technological future, NSF has initiated programs to: 1) recruit 10,000 skilled CS teachers for schools and transform the CS curriculum; 2) introduce computational thinking into STEM

education at all levels (C-STEM); and 3) create a new conceptual framework bringing together the fields of pedagogy and learning science. But this work is just beginning. A new Administration initiative coordinated by NIST is launching a multi-pronged educational effort specifically targeting cybersecurity, with the goals of increasing public awareness; expanding cybersecurity education and training at all levels; recruiting skilled cybersecurity workers for Federal missions; and boosting cybersecurity training for Federal employees. Many NITRD agencies, such as DOE/SC, NASA, NOAA, and NSF, sponsor a variety of educational activities and Web materials for schools related to their scientific missions.

*Research and education needs:* Developing the cyber-capable society will require ongoing advances in all the areas discussed. Public awareness and education reform take time, persistence, sustained coordination of efforts, and high-visibility support in every sector. It is possible that a great national challenge, similar to winning the space race in the 1960's, may also be needed to focus the attention of students and their parents, educators, and the public at large on the strategic importance of acquiring IT knowledge and skills.

**Technologies To Empower 21st Century Learning**

In the NITRD vision, IT capabilities yet to be invented will play as critical a role in education as existing IT capabilities already play in advanced scientific discovery. Today's notions of "learning technologies" will give way to concepts informed by new knowledge about the bio-chemistry of the human brain and about how the human mind develops and acquires, stores, and uses information in perceiving, thinking, and acting. Such advances in neuroscience, and parallel gains in educational and social-science research as well as in machine intelligence, will provide a basis for learning systems tailored to individuals at every stage of cognitive development.

The personalized learning system might, for example, constantly capture and update all knowledge about how individuals best learn and retain information and operationalize that knowledge in cyber tools customized for each learner. A set of cyber tools covering all domains of knowledge would be made widely available. The key aspect of these advanced cyber technologies would be their ability, through interactions with a specific individual, to adaptively recognize that person's particular characteristics as a learner and adjust the learning experience accordingly. Teachers would no longer just "stand and deliver" but could spend more time coaching and counseling individuals and leading hands-on activities.

*Where we are now:* While most U.S. schools have computers and Internet connectivity, the availability of this infrastructure for teaching varies widely by school district and type of classroom. Studies suggest that the majority of teachers do not regularly use computing capabilities in classroom lessons; they most commonly use IT to find lesson plans and teaching ideas online. The increasing number of professional organizations of teacher educators and practicing teachers interested in using technology indicates growing awareness of the IT potential, and schools of education have incorporated IT in teacher training. Nonetheless, the K-12 system, in contrast to today's professional science and engineering domains, generally lags substantially behind in making effective use of digital capabilities.

For FY 2011, NSF proposes to launch a new agencywide multidisciplinary activity – called Cyberlearning Transforming Education (CTE) – intended to catalyze the potential of advanced learning technologies to transform STEM education at all education and training levels. The

effort will support R&D both in innovative learning technologies and in learning processes themselves, within three framing themes: (1) Anytime, anywhere learning; (2) Personalized learning; and (3) (Cyber)learning for (cyber)learning.

*Research and education needs:* To realize the NITRD vision, fundamental explorations of the basis and processes of human learning should be pursued aggressively. As noted elsewhere in this plan, expanded understanding of how the mind works is a prerequisite for building smarter, more capable digital machines. In the context of education, new multidisciplinary knowledge about human cognition and development likewise will be the basis for designing revolutionary learning environments that meet the needs of individuals, whatever their location, age, socio-economic status, or educational level. These systems will become possible through technical advances in such areas as artificial intelligence, human-machine interfaces, visualization and virtual reality technologies, data analytics, compute power, networking, and distributed systems. New generations of educators appropriately equipped to work with technologies for learning will also be necessary.

# IV. Transforming Challenge Into Strategic Opportunity

***Conclusion: What We Need to Do to Get There***

## A. The Federal Role

In 1991, the legislative mandate for coordination of Federal IT R&D focused on networking and high-end computing (then the two technologies deemed vital for Federal missions). Today, advanced information technologies of many different kinds play essential and critical roles in every high-priority government mission addressing the Nation's goals and needs, and all Federal agencies rely on IT capabilities and benefit from IT research advances.

**Collaboration in Support of Mission Requirements**

The framework for multiagency coordination of Federal IT research, continuously evolving to encompass the widening range of technologies and application domains, has proven to be an invaluable asset. Collaboration among the Federal research agencies in the NITRD Program has become an intellectual and scientific imperative imposed by the diversity, complexity, interdependencies, and dynamism of contemporary IT environments. No single agency or disciplinary skill set can span more than a fraction of the IT knowledge base and investigative frontiers. Through coordination, agencies identify common mission requirements, assure focused research efforts in time-critical R&D components, and share information on addressing IT's higher-risk, higher-payoff challenges. By leveraging each other's common interests, work products, and technical expertise, NITRD agencies derive the maximum value and cost-effectiveness from Federal research investments and assure R&D coverage across the entire spectrum of technologies and domains.

Such combined efforts produce broadly applicable results that no one agency could attain on its own and that propel innovation. For example, the field of telemedicine – including two-way telepresence, remote diagnostic and haptic devices, and robotic surgical systems – emerged out of pioneering collaborative investigations and proofs-of-concept primarily funded by NIH/NLM, NSF, DoD, NASA, and the VA.

**Infrastructure for Leadership**

The multidisciplinary reach of Federal NITRD funding is extended through R&D activities in Federal laboratories, universities, research institutions, and partnerships with industry. These diversified efforts engage thousands of researchers and their students in the intellectual challenges of networking and IT and provide the principal source of education and training for the new generations of IT researchers, inventors, entrepreneurs, and technical professionals the Nation needs. In addition, the Program's emphasis on coordination and collaboration across agencies and private-sector institutions helps promote the transfer of research results and prototypes to Federal non-research agencies and out to the marketplace.

For example, the NITRD advanced-networking agencies have collaborated with private-sector partner Internet 2 in developing an infrastructure called perfSONAR that makes it possible, for the first time, to automatically measure the performance of optical networks across multiple domains. Optical networking technology is at the leading edge of next-generation networking

speeds and capacity. Accurate performance information is needed by network operators, managers, and security enforcers to support isolating and correcting network bottlenecks and anomalies; to support the configuration of network links in dynamic environments; and to detect, isolate, and correct security breaches and problems. Information on network segments in each domain is needed, requiring cooperation among the agencies and science networks maintaining these domains and network links. As a result of NITRD efforts, the perfSONAR infrastructure is being adopted by science networks, commercial network managers, and international science networks through cooperative development and deployment.

The focus of the NITRD Program on next-generation technological capabilities to support the Nation's vital interests also provides national leadership in defining research needs in critical technologies, such as high-end systems, advanced networking, cyber-physical systems, and cyber security and information assurance. In addition, the NITRD research portfolio exemplifies and promotes the multidisciplinary thinking and approaches that are increasingly necessary to develop highly complex systems of systems with myriad heterogeneous components.

## B. Partnership Synergies Are Essential

Since the dawn of the digital age, U.S. technological and economic innovation has been fueled by continuing cycles of basic exploration, experimentation, prototype implementation, feedback, and deployment of novel products. This "ecosystem" for innovation encompasses Federal researchers, private-sector researchers and developers, students and educational institutions, entrepreneurs, industries, and end users.

Today, such interactions are more necessary than ever, and many forms of engagement are needed at global scales. Strengthening security in cyberspace, for example, will require international collaboration among governments, service providers, researchers, standards organizations, law enforcement, and other stakeholders on security architectures and protocols, crime investigation, and identity management, among other issues. In a related area – the robustness and end-to-end performance of the U.S. cyber infrastructure – partnerships between government and industry in realistic-scale testbeds, technology standards, and prototype deployments of new infrastructure models such as clouds are essential to speed research advances in networking and network security to the marketplace. Innovative partnerships with the open-source software development community also should be pursued, including policy frameworks that recognize intellectual property. The exploration initiated by NITRD agencies in recent years of open-source approaches to address the challenges of complex and ultra-scale software should be continued and expanded.

History demonstrates that a sustained commitment to foundational long-term R&D is required to maintain a full pipeline of scientific findings and concepts for society's future. This role is uniquely played by Federal research investments in advanced sciences, technology, and engineering, which complement and balance industry's necessary R&D focus on bringing new products rapidly to market. NITRD's vision for the future aligns with the strategic plans of its members to pursue such fundamental technological advances as essential means of addressing the Nation's most critical needs. The NITRD Strategic Plan recognizes, however, that the scale and complexity of 21$^{st}$ century national challenges demand ongoing outreach and partnership development by the NITRD enterprise to generate synergies in the Nation's broader IT ecosystem that can accelerate beneficial advances for society and the world.

The NITRD agencies look forward to working with the Administration, the research community, the private sector, and IT stakeholders everywhere to realize the future of promise described in this strategic plan.

THE END